



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Signal-to-Noise Behavior for Matches to Gradient Direction Models of Corners in Images

David W. Paglieroni, Siddharth Manay

February 16, 2007

SPIE Defense and Security Symposium
Orlando, FL, United States
April 9, 2007 through April 13, 2007

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Signal-to-Noise Behavior for Matches to Gradient Direction Models of Corners in Images ¹

David W. Paglieroni and Siddharth Manay ²

ABSTRACT

Gradient direction models for corners of prescribed acuteness, leg length, and leg thickness are constructed by generating fields of unit vectors emanating from leg pixels that point normal to the edges. A novel FFT-based algorithm that quickly matches models of corners at all possible positions and orientations in the image to fields of gradient directions for image pixels is described. The signal strength of a corner is discussed in terms of the number of pixels along the edges of a corner in an image, while noise is characterized by the coherence of gradient directions along those edges. The detection-false alarm rate behavior of our corner detector is evaluated empirically by manually constructing maps of corner locations in typical overhead images, and then generating different ROC curves for matches to models of corners with different leg lengths and thicknesses. We then demonstrate how corners found with our detector can be used to quickly and automatically find families of polygons of arbitrary position, size and orientation in overhead images.

Keywords: Corners, Gradient Direction Matching, Signal-to-Noise Ratio, ROC Curves, Polygons

1. INTRODUCTION

Corners in image scenes represent an important class of salient points. They are used in problems that involve (1) extracting polygons associated with man-made objects (such as buildings) from overhead images [8,9], (2) establishing correspondences between points in overlapping images for image registration [22,21], stereovision, or sensor calibration [19,23], (3) detecting and tracking motion in video sequences [4,20], etc.

There are two major approaches to corner detection. In the first approach, edge curves are extracted directly from the image, or the image is segmented and region borders are treated as boundaries. Corners are then detected as curve or boundary points at which a significant change in boundary direction or curvature occurs [6,7,16]. This approach is invariant to corner acuteness and orientation, but in practice, it must be limited to images whose scene content contains mostly well-defined boundaries. Images are more likely to have this property when the acquisition conditions are tightly controlled (as in a laboratory, manufacturing, or industrial setting).

In the second approach, corners are detected by analyzing pixel neighborhoods. Algorithms of this type are not necessarily invariant to corner acuteness and orientation. However, because it does not rely on edge detection or image segmentation, this second approach is more robust than the first. It at least has the potential to handle more difficult images with scene content characterized by poorly defined boundaries, background clutter, and contrast variations. Images are more likely to have these properties when the acquisition conditions cannot be tightly controlled (as for overhead images of areas on the ground). Many methods operate over small pixel neighborhoods and are based on pixel contrast [15], gradient magnitude and direction [5,10,24] or Gaussian curvature [4]. Because these methods are so noise sensitive, the images are sometimes first smoothed, and this unfortunately reduces precision in the locations of detected corners. Techniques that use features derived from various wavelet or Gabor function transforms of the image attempt to detect pixel neighborhoods whose spatial and spectral features are consistent with corners [2,13]. Small morphological structuring elements have also been applied to gray-scale images for corner detection [11,17]. Corner locations tend to be more accurate for corners produced by the later methods because they do not require the image to first be smoothed. However, the basic difficulty with all of the methods is that since they are all sensitive to image noise, scene clutter and contrast variations, they can all potentially produce large numbers of false alarms and miss many corners.

Like other methods based on the second approach to corner detection, our method uses pixel gradient information and can be characterized as a model (or template) matching technique. However, our method should be less sensitive to

¹ This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48, UCRL-.

² Lawrence Livermore National Laboratory paglieroni1@llnl.gov, manay2@llnl.gov

noise and clutter than other methods because (1) the corner model does not depend on gradient magnitude, (2) legs in the corner model do not have to be excessively short, and (3) only pixels close to legs in the model contribute to corner matches, so background clutter is implicitly excluded. Specifically, we construct gradient direction models for corners by specifying corner acuteness, leg length, and leg thickness. The direction of the vector normal to the leg is then assigned to each leg pixel (see Section 2). As discussed in Section 3, our *Gradient Direction Matching* (GDM) method compares gradient direction models for corners to fields of pixel gradient directions using an interesting measure of similarity [14]. Although our GDM-based corner detector is not invariant to corner orientation or acuteness (and in fact uses a model for each orientation and acuteness), it is still computationally efficient because the GDM matching measure can be quickly evaluated with the Fast Fourier Transform (FFT). In Section 4, methods for dealing with ambiguous corner matches and coincident corners are developed.

In this paper, the signal strength associated with a corner match is discussed in the context of the lengths of its legs in the image. The noise associated with a corner match is discussed in the context of coherence in gradient direction for pixels along its legs in the image. We model corners as pairs of line segments (legs) emanating from a common vertex (see Section 2). Matches in the image to models of corners with short thin legs tend to have lower SNR's than matches to models of corners with longer more thick legs. By varying leg length and thickness, we are able to study the impact of signal-to-noise ratio (SNR) on corner detection performance. In Section 5, corner detection performance is assessed by generating receiver operating characteristic (ROC) curves of detection vs. false alarm probability for matches to models of corners of fixed acuteness but with various leg lengths and thicknesses. ROC curves are generated by comparing manually chosen corners (truth) to corners automatically extracted from two overhead images (one urban, one more rural) as unambiguous matches to the corner model (detections). We then demonstrate how the automatically detected corners can be used to quickly and automatically find families of polygons of arbitrary position, size and orientation in overhead images.

2. GRADIENT DIRECTION MODELS FOR CORNERS

Fig.1 depicts the model of a corner at a particular orientation as a pixel raster. In the model, the corner acuteness is $\alpha \in (0, \pi)$ (in practice, a more constrained range such as $\alpha \in [\pi/8, 7\pi/8]$ is used). Each leg has a length of $L > 0$ pixels (in practice, $L \geq$ say 3 pixels), and a thickness of $T \geq 1$ pixels. The column and row coordinates for pixels in the raster range from $-L$ to L (the raster has width $w = 2L + 1$), and for convenience, the vertex column and row coordinates are $(c, r) = (0, 0)$. Our immediate goal is to assign weights $b(c, r)$ and gradient direction angles $\beta(c, r)$ to each of the w^2 pixels.




Fig.1 Model of a corner.

As shown in Fig.1, let θ be the orientation of the corner model. Since corners in an image can occur at any orientation, each of N rotated corner models must be generated and matched separately to the image. It has been demonstrated empirically that it is sufficient to choose N close to πL (half the circumference of a circle with a radius of L pixels). By choosing the number of model orientations to be the multiple of 8 closest to πL , i.e.,

$$(1) \quad N = 8 \cdot \text{int} \left[\frac{\pi L}{8} + 0.5 \right]$$

the number of corner models whose legs occur at orientations of $k\pi/4$, $k = 0 \dots 7$ is maximized (these are the legs that map most cleanly to the pixel grid). Note that because our detector is not invariant to corner orientation, the number of corner model orientations to match, and thus the computational complexity of corner detection, increase linearly with leg length. Also, since our method is not invariant to corner acuteness, multiple models are needed for multiple corner acutenesses, although in practice, each model can be used for a range of acutenesses. The orientations of the corner models are

$$(2) \quad \theta_n = \frac{2\pi n}{N} + \frac{\alpha}{2}, \quad n = 0 \dots N-1$$

and the legs themselves have orientations of

$$(3) \quad \phi_n = \theta_n \pm \frac{\alpha}{2}$$

Suppose leg \mathcal{L} of a corner is a line segment of length L emanating from the origin at an angle ϕ . The endpoints of \mathcal{L} are $(0,0)$ and

$$(4) \quad (c', r') = (L \cos \phi, L \sin \phi)$$

The point on the line containing \mathcal{L} closest to (c, r) is

$$(5) \quad (c^*, r^*) = [c'(c'c + r'r) / L^2, r'(c'c + r'r) / L^2]$$

where

$$(6) \quad (c^*, r^*) \in \mathcal{L} \Leftrightarrow c^*(c^* - c') < 0 \text{ or } r^*(r^* - r') < 0$$

The distance from the center of pixel (c, r) to \mathcal{L} is

$$(7) \quad d([c, r], \mathcal{L}) = \begin{cases} \|[c^* - c, r^* - r]\| & [c^*, r^*] \in \mathcal{L} \\ \min(\|[c, r]\|, \|[c' - c, r' - r]\|) & \text{otherwise} \end{cases}$$

With this background, the orientations of corner model pixels are assigned as

$$(8) \quad \beta(c, r) = \phi + \frac{\pi}{2} \quad \text{if } d([c, r], \mathcal{L}) < T \text{ and } (c, r) \neq (0, 0)$$

Let P be the set of all corner model pixels assigned nonzero weights. This set is populated as follows:

$$(9) \quad (c, r) \in \text{leg } \mathcal{L} \text{ of thickness } T \text{ if } d([c, r], \mathcal{L}) < T \text{ and } (c^*, r^*) \in \mathcal{L}$$

$$(10) \quad (c, r) \in P \text{ if } (c, r) \in \text{leg } 0 \text{ or leg } 1, \text{ but not both}$$

Then, $\forall (c, r) \in P$, the weights of the corner model pixels decrease linearly from 1 to 0 as the distance to the nearest leg increases from 0 to T , and remain at 0 beyond T :

$$(11) \quad b(c, r) = \max[0, 1 - d([c, r], \mathcal{L}) / T] \in [0, 1]$$

By convention, the GDM algorithm requires these weights to be normalized so that they sum to one:

$$(12) \quad b(c,r) \leftarrow b(c,r) / \sum_{(c',r') \in P} b(c',r')$$

3. GRADIENT DIRECTION MATCHING

In an image u , gradient estimates of the form

$$(13) \quad \hat{u}(c,r) \triangleq A(c,r) e^{j\theta(c,r)} = \frac{1}{A_0} \cdot \sum_{(c',r') \in R(c,r|\rho)} [u(c',r') - u(c,r)] \cdot \frac{(c'-c) + j(r'-r)}{[(c'-c)^2 + (r'-r)^2]^{k/2}}$$

have been used to estimate pixel gradient directions [3], where $A(c,r) \geq 0$ represents gradient magnitude, $\theta(c,r) \in [-\pi, \pi]$ represents gradient direction, and $R(c,r|\rho)$ is the neighborhood containing all pixels that intersect a disk of radius ρ centered on (c,r) . Larger ρ values promote greater noise suppression whereas larger k values promote less smoothing. The fine (as opposed to coarse) gradient estimates produced by the Sobel estimator ($(\rho, k) = (1, 2)$) are suitable for corner detection ([18]).

The factor

$$(14) \quad A_0 = A_{crit} / A_{contrast}$$

is used to normalize $A(c,r)$ to units of 8-bit gray-level contrast. A_{crit} is derived from 8 bit pixel intensity contrast by computing $A(c,r)$ with A_0 set to 1 in equation (13) from a bi-level image of size $m \times m$ ($m = 2\rho + 1$) with the first ρ columns set to 0 and the last $\rho + 1$ columns set to some just perceptible contrast of $A_{contrast}$ in an 8-bit image (say $A_{contrast} \approx 12$). Then, the image pixel masking function

$$(15) \quad a(c,r) = \begin{cases} 1 & A(c,r) \geq A_{contrast} \\ 0 & \text{otherwise} \end{cases}$$

is used to mask out gradient direction estimates for low-contrast pixels (which tend to be noisy) while weighting gradient directions associated with all unmasked pixels equally.

Following [14], the Gradient Direction Matching (GDM) measure between an image u and a gradient direction model for a corner at orientation θ_n is given by

$$(16) \quad \begin{aligned} S_n(\Delta_c, \Delta_r) &= \sum_{(c,r) \in P_n} a(c+\Delta_c, r+\Delta_r) b_n(c,r) \left[2 \cos^2 [\theta(c+\Delta_c, r+\Delta_r) - \beta_n(c,r)] - 1 \right] \\ &= \sum_{(c,r) \in P_n} a(c+\Delta_c, r+\Delta_r) b_n(c,r) \cos 2[\theta(c+\Delta_c, r+\Delta_r) - \beta_n(c,r)] \in [-1, 1] \end{aligned}$$

Note that $S_n(\Delta_c, \Delta_r) = 1$ for full correlation, -1 for full anti-correlation, and close to 0 for correlation of a structured model to a random gradient direction field. Full correlation means that the model and pixel gradient directions are the same or opposite everywhere. Full anti-correlation means they are orthogonal everywhere.

Equation (16) can be efficiently evaluated using the FFT. If u has W columns and H rows ($W, H \geq w$, where $w = 2L + 1$ is the width of the square raster for the corner model), it can be shown that

$$\begin{aligned}
(17) \quad S_n(\Delta_c, \Delta_r) &= \text{Re} [B_n^*(-\Delta_c, -\Delta_r) * \Theta(\Delta_c, \Delta_r)] \\
&= \frac{I}{WH} \text{Re} \left[\text{IDFT} [\text{DFT} [B_n^*(-\Delta_c, -\Delta_r)] \cdot \text{DFT} [\Theta(\Delta_c, \Delta_r)]] \right]
\end{aligned}$$

where

$$(18) \quad B_n(c, r) \triangleq b_n(c, r) e^{j2b_n(c, r)}, \quad \Theta(c, r) \triangleq a(c, r) e^{j2\theta(c, r)}$$

$B_n^*(c, r)$ is the complex conjugate of $B_n(c, r)$, “DFT” and “IDFT” denote un-scaled forward and inverse two-sided (2D) Discrete Fourier Transforms, “*” denotes “convolution”, “.” denotes term-wise product, and the convolution property of the DFT has been used. The last expression in equation (17) is only valid for $\Delta_c = 0 \dots W - w$ and $\Delta_r = 0 \dots H - w$. If W and H are both powers of 2, then the DFT’s can be efficiently evaluated using the FFT algorithm. If not, u can first be zero-padded to force the number of rows and columns to be a power of 2.

4. CORNER AMBIGUITY AND COINCIDENCE

Corners correspond to special local maxima in the *GDM surface*

$$(19) \quad S(\Delta_c, \Delta_r) = \max_{n=0 \dots N-1} S_n(\Delta_c, \Delta_r)$$

The detection rule for corners requires the model match to exceed some critical value $S_{crit} > 0.5$ and to be an *unambiguous* local maximum within a neighborhood of radius Δ (the *disambiguation distance*):

$$(20) \quad S(c, r) \geq S_{crit} \quad \text{and} \quad S(c, r) = \max_{(c', r') \in R(c, r | \Delta)} S(c', r')$$

In our studies, we use a disambiguation distance of $\Delta = 4$ pixels. An efficient algorithm for match disambiguation by proximity is described in [14]. An image (courtesy of CASIL, [1]) and its GDM surface with respect to a corner model for which $[\alpha, L, T] = [\pi/2, 6, 1]$ are shown in Fig.2. The numbers refer to a few examples of unambiguous local maxima (right angle corner detections).

An unambiguous corner match can be represented by the state vector

$$(21) \quad \mathbf{v} = [\alpha, c, r, S_0 \dots S_{N-1}]$$

It is possible for corners with the same acuteness but different orientations to occur at the same location. Such corners are said to be *coincident*. Coincident corners can be detected by examining the sequence of S values in the state vector for an unambiguous corner match. In particular, a set of $m \geq 1$ coincident corners

$$(22) \quad \mathbf{v}_k(\mathbf{v}) = [\alpha, c, r, S_k(\mathbf{v}), \theta_k(\mathbf{v})] \quad k = 0 \dots m-1$$

can be generated from the unambiguous corner match \mathbf{v} by seeking local maxima in its sequence of S values. It is reasonable to designate the maximum number of allowable coincidences (local maxima) to be

$$(23) \quad M = \text{int} \left[\frac{2\pi}{\alpha} + 0.5 \right] > 1$$

in which case, the minimum allowable separation between successive local maxima is

$$(24) \quad \Delta_M = \max [1, \text{int} (N / M) - 1]$$

The local maxima in $\{S_0 \dots S_{N-1}\}$ can be found using a simple 3-step process: (1) Compute the set statistics $(\mu_S, \sigma_S, S_{max})$. (2) Find all elements S_m^* , $m = 0 \dots m^*-1$ from the set that are local maxima within a cyclic offset of $\pm \Delta_M$. (3) Remove all $S_m^* : S_m^* < \max [S_{crit}, \min (\mu_S + \sigma_S, S_{max})]$.




Fig.2 (a) Example overhead image of an urban area (courtesy of CASIL). (b) GDM surface for a block within the image from (a) with respect to a corner model for which $[\alpha, L, T] = [\pi/2, 6, 1]$. (c) Some selected unambiguous local maxima (right angle corner detections).

5. CORNER DETECTION PERFORMANCE

Matches in the image to models of corners with short thin legs tend to have lower SNR's than matches to models of corners with longer more thick legs. In this Section, we empirically study the impact of SNR on corner detection performance by varying the leg length and leg thickness parameters L and T for right angle corner models ($\alpha = \pi/2$).

Images of size 1024 x 1024 for an agricultural area and an urban area are shown in Fig.3 (courtesy of CASIL [1]). Examples of right angle corner detections are shown overlaid as colored dots. Red dots represent manually chosen corners (the "truth" baseline). Cyan dots represent corners automatically extracted with our algorithm using a corner model with parameters $(L, T) = (6, 2)$ and $S_{crit} = 0.75$ in equation (20). Orange dots represent overlapping red and cyan dots. Detected corners (cyan dots) are considered to be overlapping with true corners (red dots) if the distance between them does not exceed $\varepsilon = 1.5\Delta = 6$ pixels (the *detection distance*).




Fig.3 1024 x 1024 overhead images of (a) agricultural area and (b) urban area (courtesy of CASIL). Right angle corner detections for corner model parameters of $L = 6$ and $T = 2$ with $S_{crit} = 0.75$ are shown as colored dots (red dots: manual detections, cyan dots: automated detections, orange dots: overlapping manual and automated detections).

For a given choice of L and T in the corner model, a ROC curve of detection vs. false alarm probability for corner matches can be constructed by allowing S_{crit} to vary in equation (20). As S_{crit} decreases, the detection and false alarm probabilities can both be expected to increase. Coincident corners and corner orientations are not factored into our ROC curves.

For each image in Fig.3, ROC curves corresponding to various choices of L and T in the corner model are shown in Fig.4-5. In each case, the worst performance was realized by using $L = 3$ and $T = 1$ (the model corresponding to the lowest SNR detections). For $L = 3$, detection performance was much better for $T = 2$ than for $T = 1$. For $L \geq 6$, corner detection performance was comparable for $T = 1$ and $T = 2$. Detection performance did not improve, and in fact began to suffer, as T was increased beyond 2. This is to be expected because the background is more effectively excluded when a smaller value for T is used.

Detection performance also increased as L was raised from 3 to 6, but then steadily decreased as L was increased beyond 6. In each case, the best performance was realized by using a corner model for which $L = 6$ and $T = 2$, although, comparable performance was realized by using $L = 6$ or 9 and $T = 1$ or 2. Evidently, SNR is influenced not only by leg lengths and thicknesses in the corner model, but also independently by typical leg lengths for corners actually present in the image. In particular, one can expect the signal level to increase for matches to corner models as L increases towards the typical leg length for corners actually present in the image. As L increases beyond that point, one can expect the signal level to decrease for matches to corner models.

The ROC curves in Fig.4-5 suggest that even with our corner detector, it is very difficult to achieve a high probability of corner detection (say beyond 0.8) without introducing a large population of false alarms. One might thus be compelled to ask whether or not corner detection performance can be improved by combining matches to corner models with different leg lengths. A composite surface of matches can be formed by using a *composition rule* that combines surfaces of matches corresponding to models with various leg lengths such that at each pixel, the largest S value is used. Of course, the model that gives rise to the best match can vary from pixel to pixel. Unfortunately, for the images in Fig.3, we found that when this composition rule was used to combine match surfaces for $T = 2$ and $L = 3, 6, 9, 12, 15$, corner detection performance did not improve beyond what was observed for $(L, T) = (6, 2)$. Evidently, this particular composition rule promotes detections and false alarms simultaneously. Nevertheless, it may be possible to devise composition rules that tend to promote detections and demote false alarms (a topic for future research).

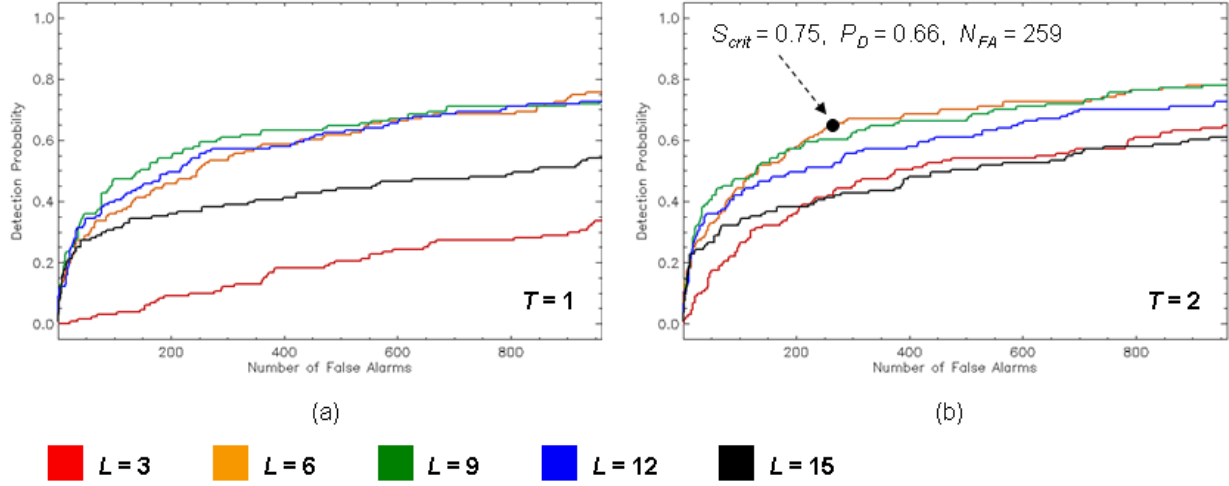


Fig.4 ROC curves for (a) $T = 1$ and (b) $T = 2$ corresponding to various choices for L in the corner model for the agricultural area image in Fig.3(a). The dot on the orange curve in Fig.4(b) ($L = 6$, $T = 2$) corresponds to the cyan corner detections in Fig.3(a).

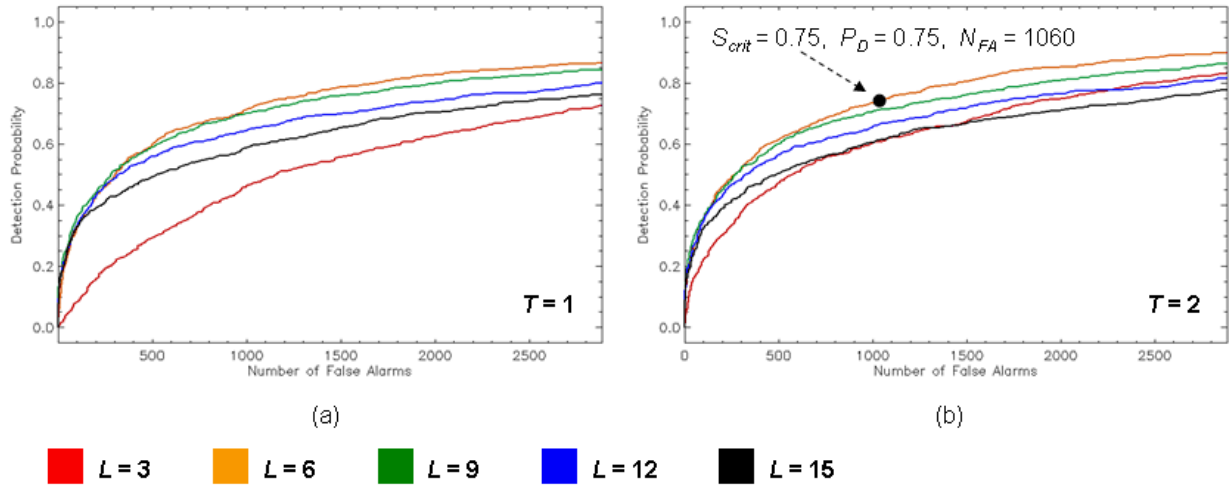


Fig.5 ROC curves for (a) $T = 1$ and (b) $T = 2$ corresponding to various choices for L in the corner model for the urban area image in Fig.3(b). The dot on the orange curve in Fig.5(b) ($L = 6$, $T = 2$) corresponds to the cyan corner detections in Fig.3(b).

Let us conclude this Section by demonstrating how corners detected using a gradient direction matching approach can be used to quickly and automatically find families of polygons of arbitrary position, size and orientation in overhead images. Fig.6 shows rectangle and T-shaped polygons extracted from the urban area overhead image in Fig.3(b). Both sets of polygons were extracted from the same field of 2048 corner detections based on matches to a gradient direction corner model with parameters $(L, T, \alpha) = (6, 2, \pi/2)$. The rectangle matches were constrained to have side lengths from 30 to 200 pixels. The T-shaped polygon matches were constrained to have side lengths from 5 to 200 pixels. The matcher uses polygon models that assign specific values of acuteness to each corner in a fixed-length sequence along the boundary. Absolute and relative lengths of sides can either be constrained or left unconstrained by the model. Both complete and partial matches are quickly found by using the model as a guide in linking corners previously extracted from images [12]. Notice that all of the T-shaped polygon matches were coincident with rectangle matches. There were

more rectangle matches because T-shaped polygons are more specific. Although polygon extraction performance can be expected to improve with corner detection performance, these examples clearly demonstrate that many if not most of the polygons can be extracted even when some corners are missed and many false corners are detected.




Fig.6 Automated (a) rectangle and (b) T-shaped polygon extraction examples from the same field of corners ($L = 6$, $T = 2$) detected automatically in the urban area image (Fig.3(b)).

6. SUMMARY AND CONCLUSIONS

A new practical method for detecting corners of designated acuteness in images has been developed. Our method analyzes pixel neighborhoods, and can be characterized as a model (or template) matching technique. Gradient direction models for corners of prescribed acuteness, leg length, and leg thickness are constructed by generating fields of unit vectors emanating from leg pixels that point normal to the edges. An FFT-based algorithm is then used to quickly match models of corners at all possible positions and orientations in the image to fields of gradient directions for image pixels. Our method has several properties that enable it to handle noise and clutter in images:

- (1) Our corner model does not depend on gradient magnitude.
- (2) Legs in the corner model do not have to be excessively short.
- (3) Only pixels close to legs in the model contribute to corner matches, so background clutter is implicitly excluded.

The signal strength associated with a corner match was discussed in the context of the lengths of its legs in the image. The noise associated with a corner match was discussed in the context of coherence in gradient direction for pixels along its legs in the image. Matches in the image to models of corners with short thin legs tend to have lower SNR's than matches to models of corners with longer more thick legs. By varying leg length and thickness, we studied the impact of signal-to-noise ratio (SNR) on corner detection performance. Not surprisingly, we found that detection performance tends to improve with SNR. However, we discovered that SNR is strongly influenced by two independent factors:

- (1) lengths of legs in the corner model
- (2) typical leg lengths for corners actually present in the image

Starting with short legs of length L in the corner model, we found that corner detection performance improves steadily as L increases (due to an increase in the SNR), but only up to a point. As L increases beyond typical leg lengths for corners actually present in the image, the SNR begins to decrease again, and detection performance grows steadily worse. For the images in our study, the optimal choice for leg lengths in the corner model was close to $L = 6$ pixels.

We also discovered that even with our corner detector, it is very difficult to achieve a high probability of corner detection without introducing a large population of false alarms. It might be possible to combine match surfaces for corner models with different leg lengths into a single composite match surface that takes into account contributions from corners at different spatial resolutions in order to provide better corner detection performance. The problem of devising composition rules that amplify detections and attenuate false alarms is an important topic for future research because any applications that rely on corners will benefit from improved corner detection capabilities. For example, one can expect polygon extraction performance to improve when better corners are supplied as input. Nevertheless, it will still be important to use corner-based polygon extractors that can cope with missing corners and large numbers of false alarms.

REFERENCES

- [1] CASIL: California Spatial Information Library. <http://archive.casil.ucdavis.edu/casil/>.
- [2] C. H. Chen, J. S. Lee and Y. N. Sun, "Wavelet transformation for Gray-Level Corner Detection", *Pattern Recognition*, Vol.28, No.6, 1995, pp.853-861.
- [3] B. Y. Chen and D. W. Paglieroni, "Using Gradients, Alignment and Proximity to Extract Curves and Connect Roads in Overhead Images", *Proc. SPIE Defense and Security Symposium*, 17-21 April 2006, Orlando (Kissimmee) FL.
- [4] L. Dreschler and H. Nagel, "Volumetric Model and 3-D Trajectory of a Moving Car Derived from Monocular TV-Frame Sequence of a Street Scene", *Proc. IJCAI*, 1981, pp.692-697.
- [5] J. Q. Fang and T. S. Huang, "A Corner Finding Algorithm for Image Analysis and Registration", *Proc. AAAI Conf.*, 1982, pp.46-49.
- [6] M. A. Fischler and R. C. Bolles, "Perceptual Organization and Curve Partitioning", *IEEE Trans. PAMI*, Vol.8, No.1, January 1986, pp.100-105.
- [7] H. Freeman and L. S. Davis, "A Corner Finding Algorithm for Chain Coded Curves", *IEEE Trans. Comput.*, Vol.26, 1977, p.297-303.
- [8] A. Huertas, "Corner Detection for Finding Buildings in Aerial Images", *USCIP Report 1050*, USC, 1981, pp.61-68.
- [9] C. O. Jaynes, F. Stolle, and R. T. Collins, "Task Driving Perceptual Organization for Extraction of Rooftop Polygons", *Proc. IEEE Workshop on Applications of Computer Vision*, December 1994, pp.152-159.
- [10] L. Kitchen and A. Rosenfeld, "Gray Level Corner Detection", *Pattern Recognition Letters*, Vol.1, 1982, pp.95-102.
- [11] R. Laganriere, "A Morphological Operator for Corner Detection", *Pattern Recognition*, Vol.31, No.11, 1998, pp.1643-1652.
- [12] S. Manay and D. W. Paglieroni, "Matching Flexible Polygons to Fields of Corners Extracted from Images", *Lawrence Livermore National Laboratory Technical Report UCRL-CONF-226354*, submitted to CVPR 2007, Minneapolis, MN.
- [13] B. S. Manjunath, C. Shekhar and R. Chellappa, "A New Approach to Image Feature Detection with Applications", *Pattern Recognition*, Vol.29, No.4, 1996, pp.627-640.
- [14] D. W. Paglieroni, W. G. Eppler and D. N. Poland, "Phase Sensitive Cueing for 3D Objects in Overhead Images", *SPIE Defense and Security Symposium: Signal Processing, Sensor Fusion and Target Recognition XIV*, *Proc. SPIE*, Vol.5809, 28-30 March 2005, Orlando FL.
- [15] K. Paler, J. Foglein, J. Illingworth and J. Kittler, "Local Ordered Grey Levels as an Aid to Corner Detection", *Pattern Recognition*, Vol.17, No.5, 1984, pp.535-543.
- [16] T. Pavlidis and S. L. Horowitz, "Segmentation of Plane Curves", *IEEE Trans. Comput.*, Vol.23, August 1974, pp.860-870.
- [17] A. Sobania and J. P. O. Evans, "Morphological Corner Detector Using Paired Triangular Structuring Elements", *Pattern Recognition*, Vol.30, No.7, July 2005, pp.1087-1098.
- [18] I. E. Sobel, "Camera Models and Machine Perception", *Ph.D. Dissertation*, Stanford University, Palo Alto, CA, 1970.

- [19] N. A. Thacker and J. E. W. Mayhew, "Optimal Combination of Stereo Camera Calibration from Stereo Images", *Image Vision Comput.*, Vol.9, 1991, pp.27-32.
- [20] P. Tissainayagam and D. Suter, "Object Tracking in Image Sequences Using Point Features", *Pattern Recognition*, Vol.38, 2005, pp.105-113.
- [21] J. Weng, N. Ahuja and T. S. Huang, "Matching Two Perspective Views", *IEEE Trans PAMI*, Vol.14, 1992, pp.806-825.
- [22] S. Yam and L. S. Davis, "Image Registration Using Generalized Hough Transform", *Proc. IEEE Conf. PRIP*, 1981, pp.526-533.
- [23] Z. Zhang, R. Deriche, O. D. Faugeras , and Q. T. Luong, "A Robust Technique for Matching Two Uncalibrated Images Through the Recovery of the Unknown Epipolar", *J. Artif. Intell.*, Vol.78, 1994, pp.87-119.
- [24] O. A. Zuniga and R. Haralick, "Corner Detection Using the Facet Model", *Proc. IEEE CVPR*, 1983, pp.30-37.